

**TOWARDS SUSTAINABLE QUALITY OF SERVICE IN
INTERCONNECTION AGREEMENTS: IMPLICATIONS FROM
INFORMATION ASYMMETRY**

Ioanna D. Constantiou

eLTRUN (eLectronic TRading Research UNit), Athens University of Economics and Business,
Department of Management Science and Technology, 47a Evelpidon & Lefkados 33, Athens, Greece
Tel.: +301 8203663, Fax: +301 8203664
ioanna@aueb.gr

Nikolaos A. Mylonopoulos

ALBA (Athens Laboratory of Business Administration), Athinas Ave & Areos Str 2A, 16671 Athens, Greece,
Tel.: +301 8964531 Fax: +301 8964737
nmylonop@alba.edu.gr

ABSTRACT

This paper analyses the structure of the Internet marketplace and the business relationships of key players involved in network services provision. A brief overview of existing pricing policies and research work in this area is presented and some new issues are introduced. We believe that the role of information asymmetry is critical when considering agreements for Internet access and interconnection. In negotiation and contract preparation, information asymmetry gives rise to adverse selection. The current structure of connectivity agreements does not address information asymmetries thus allowing the possibility of opportunistic behaviour in the form of moral hazard. Inasmuch as interconnection agreements involve sharing and/or exchanging network resources, either party will tend to exploit the agreement to its own advantage (i.e. conserving its own resources) and, possibly, to the detriment of the other (i.e. over-utilising the other's resources). The discussion focuses on interconnection agreements between Internet Service Providers, namely peering and transit. The paper concludes with an outline of an incentive compatible mechanism that can sustain quality of service requirements in interconnection agreements.

1. INTRODUCTION

The Internet is structured hierarchically, comprising three main levels of participants, namely, end users, Internet Service Providers (ISP) and Internet Backbone Providers (IBP). End users are at the bottom of the hierarchy and access the Internet via ISPs. End users include residential and business customers. At the top of the hierarchy, IBPs own high speed and high capacity networks to provide global access and interconnectivity. They sell primarily wholesale Internet connectivity services to ISPs [Shriganesh, 1997]. ISPs then resell connectivity services or add value and sell new services to their customers. However, IBPs may also get involved in ISP business activities by selling retail Internet connectivity services to

end-users. Both IBPs and ISPs provide complementary inputs to the bundled network services that end-users consume [Foros and Kind, 2000].

This hierarchical value chain for Internet connectivity involves two main kinds of pricing contracts. The first involves pricing between end-user and ISP for primary Internet access and the second involves pricing between ISP and IBP for interconnection. Thus, two markets for Internet connectivity are identified, wholesale and retail, for global access and for connectivity to end-users respectively [Huston, 1999]. In the early days when the Internet was an activity restricted to the public sector, mainly for research and education purposes, access and interconnection were public goods and their provision was organized outside competitive markets.

Internet provision and use today is primarily commercial, yet its basic architectures remain unchanged. Internet connectivity in itself possesses public good properties, the most pervasive being network externalities. The value for each individual participant, derived from the 'network of networks' increases exponentially with broader reach and greater participation. Similarly, inefficient utilization of network resources by one participant has detrimental effects on the quality of service received by others. Externalities generate powerful incentives for interconnection while setting the stage for potential opportunistic exploitation of shared network resources.

Another key characteristic of the Internet is its variety and heterogeneity. Diverse technologies, applications and services inter-operate almost seamlessly, posing heterogeneous requirements for network resources. Heterogeneous end-users have diverse expectations from the network and make use of various technologies, applications and services. It is quite remarkable that such a complex commercial and technological ecosystem seems to operate without major breakdowns and the whole world effectively relies on it for a wide range of social and economic activities.

Having said all that, there are credible signs that the simple market mechanisms governing internet connectivity are cracking under current pressures and may not be able to sustain future growth. In particular, flat rate pricing for primary access is consistent with the cost structure of Internet service provision [Mackie-Mason and Varian, 1995] but does little to control resource allocation under conditions of high demand (congestion) and/or in the presence of differentiated user demands and willingness to pay. For example a business user under pressure to complete a certain task may be willing to pay more in order to achieve higher transfer rates. Current Internet architectures mainly support 'best effort' service and pricing schemes cannot discriminate between users requiring high quality of service and casual web surfers. This problem is exacerbated during periods of high traffic or congestion, when efficient resource allocation becomes an even more pressing concern.

The problem is more complicated once we realize that even if we devise an efficient and practical mechanism for resource allocation and price discrimination for primary access [Cremer et al. 1999], actual performance depends on the conditions and behavior of several networks that mediate data transmission throughout the world. In the past, ISPs have been agreeing to service each other's traffic without charge, for their obvious mutual benefit. However, competitive market dynamics have tilted the balance in such peering agreements when, for example, one partner makes heavier use of another's resources. Commercial wholesale contracts, on the other hand, cannot always verify or enforce the agreed performance levels. For example, a wholesale network provider may disguise his low effort (e.g. neglect to upgrade bottleneck network components) as adverse system-wide demand conditions. An even more elementary problem is to agree upon what constitutes performance, effort and cost and how that is built into an effective pricing scheme [Kende and Oxman, 1999].

Information asymmetry is a key component of the problems sketched above. Participants in retail and wholesale markets for Internet connectivity lack full information regarding each other's capacity, demand, resource allocation, effort and cost. [Cukier, 1998] As a result, they cannot enforce any contracts based on performance (quality of service) and they have an incentive to act opportunistically against each other (to take advantage of the other party's poor information and deviate from agreed performance).

This paper explores current practices and research work on access and interconnection agreements and outlines future research directions for taking information asymmetry explicitly into account. It is structured as follows. The next section reviews the literature on pricing retail Internet access in the presence of quality of service (QoS) requirements. It explores the challenges faced, alternative solutions and their limitations. We argue that even if these problems are resolved at the retail end, QoS cannot be guaranteed unless the same requirements are propagated throughout the interconnectivity chain. Section 3 turns to interconnection agreements between network providers. Current practices are discussed and their weaknesses are identified. Section 4 explores in detail the nature, manifestations and implications of asymmetric information in interconnection agreements. This section also sets out the requirements for sustainable QoS expectations from such agreements. Before concluding the paper, section 5 outlines a modeling approach to characterizing an incentive compatible mechanism for interconnection, satisfying such requirements.

2. PRICING RETAIL INTERNET ACCESS

Much of the existing work on Internet access pricing adopts the view that prices should be used to achieve efficient resource allocation. In this context, the proper objective function is some measure of user satisfaction rather than cost parameters given *ex-ante*. Clearly, Internet access provision exhibits very low marginal cost, much like any other information good. Therefore, it is not surprising that competition drives actual prices to very low levels, even to nominal zero values [Jew and Nicholls, 1999] and that internet access is often bundled with other information services (e.g. AOL) [Bakos and Brynjlofsson, 1997]. Consequently, pricing turns to consumer valuations as the basis for determining how much to charge. In the presence of heterogeneous user preferences regarding quality of service (i.e. delay and packet loss), the problem becomes a standard exercise in price discrimination. Using such models it is straightforward to show that properly computed short run prices could give information about the value of capacity, and provide useful indications for network resource allocation. [Shriganesh, 1997]

Further, an important element in some models is that users should be charged in a way that reflects the negative externalities they impose on others. Hence, if the capacity of the network is constrained, a user should be charged for the fact that the packets he transmits increase congestion and therefore decrease the utility of other users. If negative externalities are not charged for, information asymmetry (for example, the fact that the ISP cannot know whether the user is a casual web surfer or an urgent business client) gives end users an incentive to act opportunistically by ‘wasting’ network resources.

The Internet raises specific problems once one turns to the much more complex issue of implementing these pricing schemes. One pricing problem is that a user would be made responsible for costs over which he has very little control as they depend on total network capacity and the behavior of other users. There are two issues here.

One is incomplete information about network capacity and demand patterns. Even if a user makes the effort to collect relevant information from research reports, reputation or past experience, his information will always be incomplete *vis-à-vis* the informed ISP who has perfect knowledge of both (even if demand is not perfectly predictable). This information asymmetry can instigate ISP opportunism in the form of moral hazard. For example, an ISP might oversubscribe its network in order to maximize profitability at the detriment of customer service. Alternatively, with a pricing mechanism that charges for negative externality, and given that actual demand is stochastic and unobservable by the user, the ISP might discriminate his charges against certain customers in favor of more profitable ones.

The second issue relates to a degree of risk that the user may have to bear under such pricing mechanisms. This risk involves receiving lower service quality for any given price or vice versa. Normally, in a principal-agent setting we expect the risk-neutral party (in this case the ISP) to absorb the risk facing the risk-averse party (in this case the end user). Inasmuch as the ISP also faces some risk from uncertain demand, this issue is further confounded and the efficient outcome is not obvious.

The second pricing problem stems from the fact that the shadow prices of capacity vary over time, and it is possible that in some parts of the network where capacity is specially limited and demand especially bursty,

these shadow prices may vary very fast. This makes it very difficult to organize a pricing scheme that tracks the correct shadow prices. Given that demand is stochastic and highly bursty, this problem can be pervasive. From a practical perspective, the user cannot handle the overhead of adjusting his behavior to dynamic, real time price variation.

In their proposal for 'smart markets', McKie-Mason and Varian [McKie-Mason and Varian, 1995, 1996, 1997], suggest the use of Vickrey auctions. In their implementation, each packet would carry a maximum price that the sender is willing to pay for the service. The network would accept to forward packets carrying a willingness to pay superior to a threshold computed in such way that the total number of packets that are transmitted equals available capacity. Although this approach yields a theoretical first best outcome, it faces obvious implementation obstacles (including, among others, network overheads and incompletely informed users). Furthermore, one might raise questions regarding the equity and fairness of the method, since some users may get served too late or never.

In a series of papers, Gupta, Stahl and Whinston [e.g. Gupta et al., 1997 and 2000] have proposed alternative approaches for pricing Internet access dynamically. At every point of time the network is monitored for congestion. The prices charged for the nodes at which congestion is severe are increased, whereas the prices for nodes at which congestion is less severe are decreased. Each user of the network is informed dynamically of the prices and can decide whether or not to send packets accordingly. The authors have carried out simulated experiments to demonstrate that under broad conditions their algorithm tracks the equilibrium prices well.

A key feature of the next generation Internet Protocol (IPv6) is that it will support priority classes that data packets will be assigned, in order to satisfy QoS requirements of the application (e.g. email vs. videoconferencing) or of the user. This feature has been exploited in much of the research reviewed here. In particular, Clark's [Clark, 1995] basic proposal gives users the opportunity to buy 'priority flags', which can be attached to especially important packets. Priority classes can be used to implement a pricing mechanism based on expected capacity. The main advantages are performance predictability for the user (even though QoS may not be entirely guaranteed) and ease of implementation for the provider (no tracking of data). Moreover, this analysis captures directly the fact that the marginal cost of traffic flow is non-zero only during congestion and that prices remain flat during non-congested periods.

Dynamic pricing mechanisms are 'spot market' proposals. Although they have significant efficiency traits [Edell and Varaiya, 1999], they have three main disadvantages. First, they are not suitable for applications requiring continuous availability of bandwidth, where demand cannot adjust dynamically. Second, they may pose significant information overheads to the network itself, which has to monitor traffic volumes on every node and notify end users. Third, they raise informational overhead for the user and are disruptive to the main service that the user seeks online.

Psychological experiments [Bouch and Sasse, 1999] have demonstrated the latter effect and elaborated on its implications. In particular, it has been shown that users need predictable quality of service and real time feedback for tasks demanding high performance, while more 'casual' or less urgent tasks do not demand continuous feedback on network performance. Moreover, user willingness to pay is positively correlated with predictability of and confidence in network performance. In other words, Internet access pricing should not only reflect dynamic (spot) network conditions but also the overall, or long term quality of service, which shapes subjective user expectations. Therefore, dynamic pricing mechanisms are limited not only because they impose information overheads to the network and to the user but also because in the long run they turn out to be less efficient than what current models anticipate.

This section discusses some approaches for service performance expectations of internet users. All these approaches face a fundamental limitation, namely that solving the problem at the retail end is insufficient unless the same requirements are propagated throughout the internet connectivity value chain. We now turn the discussion to this matter.

3. PRICING WHOLESALE INTERNET ACCESS: INTERCONNECTION AGREEMENTS

Interconnection is what makes the Internet “the network of networks”. The dominant economic driving force for interconnection between network providers is positive network externalities. Externalities result from connectivity [Baake and Wichmann, 1999], the possibility of every party connected to the Internet to be able to communicate with any other party, and from universal access, the possibility to have access to all network resources independently of the user’s physical location. Additionally, interoperability of heterogeneous technologies and user applications requires extensive connectivity. Further, dense interconnection facilitates packet routing through short paths and decreases the possibility of marginal loss, thus supporting the provision of Quality of Service (QoS).

ISPs and IBPs aiming at exploiting the benefits of interconnection have been implementing two types of agreements, namely peering and transit. As a result, the exchange of Internet traffic operates with two parallel systems [Cukier, 1998]. Peering agreements involve the exchange of traffic between the users of two networks free of charge. When ISP A peers with ISP B, traffic originates in A’s (B’s) network and terminate in B’s (A’s) network. Transit payment agreements occur when a provider wants to reach customers of some third party that he doesn't peer with. In this case he enters into transit agreement with another intermediary provider who is, in turn, interconnected to that part of the Internet. Neither peering or transit agreements guarantee quality of service.

A peering agreement between two ISPs involves exchange occurring at public and private Internet exchange points. Partners only exchange traffic between them, at the exchange point nearest to origination and termination of transfer, on a settlement-free basis also known as sender-keeps-all. The only direct costs involved are the purchase of equipment and the provision of transmission capacity needed for each partner to meet the requirements deriving from peering. Figure 1 presents a simple representation of the main connectivity arrangements described above.




Figure 1: Interconnectivity on the Internet

Connectivity providers face conflicting incentives. On one hand, they have an incentive to cooperate with one another in order to provide their customers with access to the full range of Internet users and content. On the other hand, they have an incentive to compete with one another for both retail and wholesale customers. The strategies and growth of individual networks vary significantly and this has led to some breakdowns in the peering system. [Frieden, 1998] The exponential growth of the Internet has put enormous pressure on the backbones and on the interconnection points connecting the backbones. [Cremer et al., 1999]. As a result, performance is hampered at these points and peering often turns out to be inferior in terms of service quality.

The alternative to peering agreements is transit payment agreements. There are two main differences between peering and transit. First, one partner pays another partner for interconnection and therefore becomes a wholesale customer. The partner selling transit services will route traffic from the transit customer to its own peering partners as well as to other customers [OECD, 1998]. Second, transit does not involve the same service as peering and, therefore, refusing peering in favour of transit is not a means of charging for a service that was otherwise provided free of charge. When regional ISPs pay for transit they benefit from the infrastructure investments of national or global backbones without themselves having to make the same investments. Transit gives an ISP access to the entire Internet, not just the customers of the peering partner, thus the transit provider must either maintain peering arrangements with a number of other backbones or must pay for transit from another backbone.

Many ISPs have adopted a hybrid approach to interconnection, peering with a number of ISPs and paying for transit from one or more backbones in order to have access to those backbones they do not peer with [Kende and Oxman, 1999]. Interconnection agreements are also influenced by the dynamic nature of the Internet, which often leads to a form of arbitrage that is played behind the scenes by the different ISPs negotiating new interconnection agreements. For example an IBP that provides connectivity to smaller ISPs must also interconnect with other IBPs and act similarly to foreign exchange arbiters, as he seeks to extract revenue in both directions. The resultant business environment is one characterised by a degree of fluidity. Many network providers operate both as a client and as a provider [Huston 1999].

4. IMPLICATIONS FROM INFORMATION ASYMMETRY IN INTERNET CONNECTIVITY MARKETS

As indicated in earlier sections, bilateral transactions in the market for Internet connectivity (wholesale or retail) are characterised by severe information asymmetries [Macho-Stadler and Perez-Castrillo, 1997]. Internet providers control all the information pertaining to the characteristics of their networks (e.g. capacity, usage etc.) and may or may not disclose it to potential interconnection partners. From an economic perspective, such information is critical for the structure and efficiency of interconnection agreements. Current practices are often based on the subjective perceptions of the parties involved and may not be optimal or sustainable because of asymmetrically available information.

Asymmetric information in current types of interconnection agreements gives rise to opportunistic behaviour in different guises. The first is called “backbone free riding”. A national ISP has to build and maintain a nation-wide network, connecting different regions, whereas a local ISP, concentrating on a single region does not. If both ISPs agree to interconnect, the local ISP may use national ISP capacity to service traffic between customers in distant regions. For example, when a customer of the regional ISP requests a web page from a customer of the national backbone whose server is far away, the request will be carried through the national ISP, from one region to the other and the response back. The national ISP may thus refuse to peer on the grounds that it is bearing the expense for a national infrastructure that the regional ISP can exploit at no cost. As a result, a number of ISPs include in their publicly stated peering policies that potential peer partners should be willing and able to peer at a number of geographically dispersed locations [Ergas, 2000].

The second manifestation of opportunism is called “business stealing effect”. Interconnection naturally lowers end user switching costs. End customers may switch network providers seeking better price/performance ratios without losing connectivity or access to shared network resources. Lower switching costs increase competition and, as a result, weaken ISP incentives to interconnect [Shapiro and Varian,

1998]. An alternative strategy is to raise switching cost by differentiation. An ISP may bundle exclusive services or content to its main Internet access offering in order to achieve customer lock-in (e.g. AOL).

Another example of perceived free riding that may arise in a peering relationship derives from the business strategy of an ISP. One ISP may choose for a variety of reasons to focus on providing service to users that generate high traffic volumes and use extensively the web servers of the peer ISP. In such cases the second ISP will carry extra traffic volume that will negatively affect its network performance, and decrease the quality of services provided to its own customers. If usage patterns are not reciprocal, peering is not sustainable. Opportunistic behaviour may also arise in transit agreements. When an ISP signs a transit agreement he is expecting to have global access to the Internet. It is, however, difficult to know the network coverage of his provider and the performance levels of its network. As ISPs are trying to increase their revenues through higher utilisation of their network, they often oversubscribe it. This behaviour in combination with best effort service provision may end up to increased delays and packet losses for client traffic. Thus, ISPs entering transit agreements do not always receive their expected benefits.

When ISP A is not able to identify the type of ISP B, with respect to certain characteristics that will affect the outcome of an interconnection agreement, there is an adverse selection problem. The result might be that desirable interconnections may not be agreed or that agreements may be settled under unfair or inefficient conditions. The main information components that may be asymmetrically available to candidate interconnection partners include, among others, the following.

- The types of customers. Customer demand is notoriously unsystematic and difficult to predict. However, an ISP can obtain demographic and usage characteristics (as indicators of demand patterns) of its client base. Such information is not available to third parties.
- The volume of traffic exchanged. This information is directly related to customer demand, which is not predictable. However it is possible to simulate or estimate statistically demand patterns on the basis of historical data [e.g. Gupta et al 2000].
- Presence at peering points, other peering agreements and network management. Such information concerns the business strategy of the ISP and its core competence. An ISP has no incentive to reveal this type of information that will directly reveal the cost of managing its network.
- Available capacity and resource allocation. This information includes decisions on statistical multiplexing, overbooking, attracting new customers. Resource allocation has strong implications for network performance.

Such information is critical during negotiations for peering or transit agreements. However, it is not readily available and ISPs have little incentive to reveal it or report it truthfully. Current market practices address this problem only in part. Large ISPs exert their bargaining power to extract such information from smaller potential partners. The requirements and terms of such agreements are privately communicated and undisclosed.

Information asymmetries are also manifest in the form of moral hazard after an interconnection agreement is entered into. When ISP A is not able to observe or monitor the behaviour of ISP B after an interconnection agreement, ISP B may alter its behaviour opportunistically for its own private benefit and to the detriment of ISP A (or vice versa). Moral hazard arises as a result of actions such as the following.

- An ISP may not keep upgrading his network capacity after an interconnection agreement. This will result in poorer servicing of the partner's traffic. As interconnection agreements currently are based on best effort services, such behaviour cannot be verified.
- An ISP may actively discriminate against IP packets that enter into his network from the interconnected partner when its network has high traffic. In the context of best effort services it is almost impossible to detect and verify such behaviour.

- An ISP may overbook its network in order to maximise economies of scale. To avoid congestion the ISP may delay or not admit interconnected traffic. This not the predictable outcome under ‘naturally’ arising congestion but the result of intentional unilateral overbooking.

Moral hazard appears because one ISP’s profit maximisation strategy may not be aligned with the interests of its interconnection partners and because he can hide or disguise his behaviour. The result is inefficient and unstable agreements. Incentive compatible contracts can be devised so as to safeguard interconnection agreements from opportunism and sustain the undeniable benefits of network externalities.

5. PROPOSED RESEARCH DIRECTIONS

Both peering and transit agreements often do not provide sufficient economic incentives for partners to collaborate on exploiting positive network externalities. New internet applications appear to be increasingly demanding in terms of specific network performance guarantees. We argue that new types of interconnection agreements based on contracts with incentive mechanisms, will mitigate the adverse implications of asymmetric information and will provide a sound basis for sustaining quality of service requirements.

There are two main issues open to future research on interconnection agreements.

- How can an interconnection customer ensure receiving fair treatment in a best effort service network?
- What is the appropriate pricing scheme that will induce an interconnection provider to treat client traffic with more than best effort services?

In order to apply the asymmetric information framework to the interconnection market, we identify the basic parameters in the Internet context. These parameters are effort, outcome and cost.

The effort in this context can be defined as the ISP’s decision on how to treat the incoming traffic of a customer. When customer traffic enters the provider network a decision is taken on the path it will follow within it. This decision affects the quality of interconnection with respect to average delay and packet loss rate. It is not observable by the customer, and the provider has no incentive to reveal it. The inability to verify ISP effort can be alleviated by devising pricing mechanisms that provide suitable incentives to the ISP to exert such effort as to ensure the expected performance. In effect, such mechanisms make the ISP responsible for the effort he exerts by tying his payment to the outcome after accounting for uncertain conditions.

We can assume that effort is defined in terms of the multiplexing algorithms applied by the ISP. Multiplexing algorithms can be manipulated to give different priorities to different kinds of packets according to subjective criteria. Such criteria may include, among others, the type of application being serviced (e.g. email vs. videoconferencing), the identity of the sender (or recipient) or the revenue generated by the traffic flow.

The **outcome** in the internet context can be defined by performance indicators such as average delay or packet loss rates that are observable and provide quantitative measurements of interconnection quality.

The **cost** of interconnection can be defined by taking into account network management criteria adopted and the multiplexing algorithms, selected by the ISP. An obvious definition of this cost is the opportunity cost of not serving (or reducing the quality of service for) other customers. An alternative but equivalent definition of this cost is in terms of negative externality (congestion) imposed on the network and its users. It is quite difficult to estimate this cost as it depends on parameters that an ISP may not reveal. A key parameter is the dynamic condition of the network, defined as the traffic load that is already in the network. This information is available to the provider before deciding how to treat incoming traffic. In this setting, the cost of effort is zero under some threshold traffic level and increases exponentially above that threshold. In turn, this threshold depends on total available capacity and on the multiplexing algorithm.

To demonstrate a representative example, we assume that there are two networks A and B. Network B sells backbone connectivity to A via a transit agreement. Further, we assume that a customer of network A wants to set up videoconference with another party at a distant network. Videoconference packets will pass through

network B. The application will generate revenue mainly to ISP_A, who owns the customer who, in turn, will be willing to pay more in order to ensure the desired frame rate. Network B has a key role in service delivery. According to current types of interconnection agreements, ISP_B will get no extra revenue and has no incentive to provide better performance. How can ISP_A induce ISP_B to provide performance guarantees or at least priority to its videoconference packets?

We consider a simple case with two effort levels. ISP_B maintains two virtual links of equal capacity. The average arrival rate is higher on the first link. Therefore the average delay rate will also be higher for packets forwarded through the first link. The second link is reserved for priority packets, has lower arrival rate and sustains lower delays. ISP_A would like to induce ISP_B to pass his packets through the priority link. However, ISP_A is not able to monitor or verify ISP_B effort. ISP_A will only observe the final outcome (delay or frame rate), which is expected to be the weighted average of the two different delay rates. The cost of ISP_B depends on current traffic and on the extra delay that ISP_A traffic will create on the specific link. This implies that there is an opportunity cost for carrying ISP_A traffic through the priority link. Both ISPs are maximising their profits. ISP_A prefers his traffic to go through the priority link whilst ISP_B generally prefers to route traffic through the low priority link that has lower opportunity cost. At this point the analysis is more straightforward from the point of view of ISP_B. The latter has to segment the market with such a price that only customers, who really value priority service, receive it. The price must be tied to observed frame rates while being higher than ISP_B's opportunity cost. The reality the problem is rather more complex because ISP_B's opportunity cost depends on the total demand for the priority link at the given moment in time. In a comprehensive model, the uncertainty regarding system-wide demand conditions and the risk preferences of the parties must also be taken into account.

This is a sketch of a model that addresses the question of sustaining efficient interconnection agreements. It suggests that mechanisms similar to those studied for retail Internet access can be applied to wholesale markets. The aim is to propagate sufficient incentives throughout the connectivity value chain in order to deliver performance expectations to end-users. Clearly, formal modelling is warranted to explore the properties and behaviour of such mechanisms. Furthermore, multiple provider settings need to be modelled before broader conclusions can be drawn.

6. CONCLUSIONS

This paper analyses both retail and wholesale markets for internet connectivity, giving particular emphasis on the latter. After a brief overview of extant pricing policies and related research we believe that the role of information is critical when considering bilateral business relations in the Internet marketplace. Asymmetric information in current types of interconnection agreements gives rise to opportunistic behaviour with negative implications. Interconnection agreements exhibit great instability because of the difficulties in enforcing them. Both peering and transit agreements do not provide sufficient incentives for partners to collaborate on exploiting positive network externalities. Our research approach is aimed at addressing adverse selection and moral hazard in interconnection agreements directly, with a view to devising incentive mechanisms suitable for stable interconnection agreements sustaining specific performance expectations.

REFERENCES

- Baake, P. and T. Wichmann (1999). On the Economics of Internet Peering. *Netnomics*, 1, 89-105. Efige
- Bakos, Y. and E. Brynjolfsson (1997). Bundling Information Goods: Pricing, Profits and Efficiency. In *Proceedings of the Conference on Economics of Digital Information and Intellectual Property*, Harvard University.
- Bouch, A. and M.A. Sasse (1999). It ain't what you charge, it's the way that you do it: A user perspective of network QoS and pricing". In *Proceedings of the IM'99*.
- Bouch, A. and M.A. Sasse (2000). The case for predictable network service. In *Proceedings of the MMCN'2000*. San Jose, CA, 188-195.

- Clark, D. (1995). A Model For Cost Allocation and Pricing in the Internet. In L. McKnight and J. Bailey. Internet Economics. 1st Edition. MIT Press, Cambridge, 155-168.
- Cremer, J, Rey, P. and J. Tirole (1999). Connectivity in the Commercial Internet. Institut d'Economie Industrielle (IDEI), www.idei.asso.fr/English/EPresent/index.html. Access on 14/11/00.
- Cremer, J. and C. Hariton (1999). The pricing of Critical applications in the Internet. Institut d'Economie Industrielle (IDEI), <http://www.idei.asso.fr/English/EPresent/index.html> Access on 14/11/00.
- Cukier, K.N. (1998) Peering and Fearing: ISP Interconnection and Regulatory Issues. Harvard Information Infrastructure Project, <http://www.ksg.harvard.edu/iip/iicompol/Papers/Cukier.html> Access on 14/11/00.
- Edell, R. and P. Varaiya (1999). Providing Internet access: What we learn from INDEX. In Proceedings of the INFOCOM '99.
- Ergas, H., (2000). Internet Peering: A Case Study of the ACCC's Use of its Powers Under Part XIB of the Trade Practices Act, 1974. In Proceedings of the 3rd Berlin Internet Economics Workshop.
- Frieden, R., (2000). When Internet Peers Become Customers: The Consequences of Settlement-based Interconnection. In Proceedings of the 3rd Berlin Internet Economics Workshop.
- Frieden, R. (1998). Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization, Virginia Journal of Law and Technology, 3 (Fall).
- Foros, O. and H.J. Kind (2000). National and Global Regulation on the Market of Internet Connectivity. In Proceedings of the 3rd Berlin Internet Economics Workshop.
- Gupta, A., Jukic, B., Stahl, D., and A. Whinston (1997). A stochastic equilibrium model of Internet pricing. Journal of Economic Dynamics and Control, 21.
- Gupta, A., Stahl, D., And A. WHINSTON (1999). The Economics of Network Management, Communications of the ACM, 42 (5), 57-63.
- Gupta, A., Jukic, B., Stahl, D., and A. Whinston (2000). Extracting Consumer Private Information for Implementing Incentive-Compatible Internet Traffic Pricing. Journal of Management Information Systems, 17 (1), 9-29.
- Huston, G. (1999). Interconnection, Peering and Settlements. In Proceedings of the INET99 Internet Society Conference.
- Jew, B. and R. Nicholls (1999) Internet Connectivity: Open Competition In the Face of Commercial Expansion, In Proceedings of the Pacific Telecommunications Conference.
- Kende, M. & J. Oxman (1999). The Information Interchange: Interconnection on the Internet, In Proceedings of the Common Telecommunications Research Policy Conference.
- Macho-Stadler, I. and D. Perez Castrillo (1997). An Introduction to the Economics of Information: Incentives and Contracts. 1st Edition, Oxford University Press, Oxford.
- Mackie-Mason, J.K. and H. Varian (1997). Economic FAQ on the Internet. <http://www-personal.umich.edu/~jmm/papers/FAQs/FAQs.html> Access on 14/11/00.
- Mackie-Mason, J.K. and H. Varian (1996). Some Economics of the Internet. In Sichel B. and Alexander L. (eds) Networks infrastructure, and new task for regulation, Ann Arbor: The University of Michigan Press, 107-136.
- Mackie-Mason, J.K. and H. Varian. (1995). Pricing the Internet. In Kahin B. and Keller J.(eds) Public access to the Internet, MIT Press, Cambridge MA, 269-314.
- Shapiro, C. and H. Varian (1998) Information Rules: A strategic guide to the network economy. Harvard Business School Press.
- Sriganesh, P. (1997). Internet cost structures and interconnection agreements. In L. McKnight and J. Bailey. Internet Economics. 1st Edition. MIT Press, Cambridge, MA.
- Working Party on Telecommunications and Information Services Policies, (1998). Internet Traffic Exchange: Developments and Policy. OECD.